Economics and Philosophy, **23** (2007) 373–384 Copyright © Cambridge University Press
doi:10.1017/S0266267107001563

WEIGHING LIVES – AN APPLIED ECONOMIST'S PERSPECTIVE*

Michael W. Jones-Lee
University of Newcastle Upon Tyne

1. THE VIEW FROM MOUNT OLYMPUS

Without doubt, *Weighing Lives*, like its precursor, *Weighing Goods*, is an excellent and thought-provoking piece of work. In the first place, it addresses a question of the most fundamental importance, namely: how should we aggregate the well-being of past, present and future members of the human race under the various possible states of the world that may, in the event, prevail? This involves, amongst other things, dealing with questions of aggregation across time, people and different states of the world; the issue of what constitutes a "neutral" state of well-being (i.e., one neither better nor worse than not being alive) and what is, in fact, the value of continuing to survive.

All of these and related issues are considered by its author in an incisive, systematic and clearly argued manner and I, for one, have little if any quarrel with what he has to say.

So, if we had a satisfactory means of measuring the extent of a person's well-being under given circumstances at a particular point of time, would John Broome's book help us in our efforts to make the world a better place for current and future generations to live in? The answer is probably "yes". But unfortunately, with respect to the question of whether we do in fact have a satisfactory means of measuring overall individual well-being, the answer – at least at present – appears to be an unequivocal "no".

* This paper is part of the Book Symposium on John Broome's *Weighing Lives*, published in the March 2007 issue of *Economics and Philosophy*. Due to an editorial oversight, it was not included in that issue. It is followed by the relevant part of John Broome's response, reprinted from the original symposium.

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

2. MUDDLING THROUGH IN MUDDY BOOTS

In the light of all this how, then, is one to proceed when approached by a public or private sector decision maker with a request for advice on whether or not to go ahead with an investment project or regulatory decision that can be expected, amongst other effects, to have some (usually, but not invariably) beneficial effect upon the health and/or safety of human life? Does one, for example, point to John Broome's work and say when we *do* have a satisfactory means of measuring overall individual well-being then *Weighing Lives* will provide all these answers, but until then we'll just have to toss a coin. Or does one try to do the best one can within the limits of current knowledge and capabilities. As a muddy-booted applied economist my response is to try to do my best within the current limitations.

So what *do* I do? In the first place I accept that I have no satisfactory means of measuring the extent of an individual's overall well-being and so haven't a clue how one might shift Joe Soap or Jemima Puddleduck from their current state of affairs to some idealized "optimal" condition. What I believe I can do, however, is to make a stab at saying whether a *marginal* adjustment in their current condition – treating most other aspects of their lifestyle as parametric – is a good or a bad thing. More specifically, if all other determinants of their own level of well-being are more or less unaffected, then I am prepared to regard a marginal increase in their command over marketed goods and services and/or an increase in their survival probability for a forthcoming period as making a positive contribution to their own quality of life. Thus, denoting individual i's command over marketed goods and services (or in an economist's terms, their level of "wealth") by w_i, their survival probability for a forthcoming period by p_i and all other determinants of their level of well-being by g_i, I would be perfectly happy to regard their *overall* level of well-being or (again in an economist's terminology) utility, u_i, as being represented by a well-behaved cardinal function

$$(1)\quad u_i = u_i(w_i,\ p_i,\ g_i)$$

which is increasing in all of its arguments and, for the sake of simplicity, at least twice-differentiable.

In turn, in the spirit of John Harsanyi (1955) and John Broome, I would be prepared to aggregate individual well-being into a measure of overall social welfare (or in Broome's terms, "value") by a weighted additive function

$$(2)\quad W = \sum_{i=1}^{n} a_i\, u_i$$

where the a_i are distributional weights.

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

It then follows from equations (1) and (2) that the impact on social welfare, δ W, of a project or programme entailing marginal variations δp_1, δp_2,. . .δp_n in individual survival probability for a forthcoming period will be given by

$$(3)\quad \delta W = \sum_{i=1}^{n} a_i \frac{\partial u_i}{\partial p_i} \delta p_i$$

or, equivalently, by

$$(4)\quad \delta W = \sum_{i=1}^{n} a_i \frac{\partial u_i}{\partial w_i} \left(\frac{\partial u_i}{\partial p_i} \Big/ \frac{\partial u_i}{\partial w_i} \right) \delta p_i.$$

As far as equation (4) is concerned, two things should be stressed.

- The term $a_i \frac{\partial u_i}{\partial w_i}$ represents the weight applied by the social welfare function to a marginal variation in the i^{th} individual's wealth i.e., in an economist's terms, it is the "marginal social welfare of the i^{th} individual's wealth" or, more succinctly, the "distributional weight for the i^{th} individual's wealth" and
- The term $\frac{\partial u_i}{\partial p_i} / \frac{\partial u_i}{\partial w_i}$ is the rate at which the i^{th} individual is willing to trade off own wealth against own survival probability "at the margin" i.e., in an economist's terminology, his/her "marginal rate of substitution (MRS) of wealth for survival probability."

With appropriate choice of units for a_i, equation (4) therefore provides a direct measure of the "value" of a safety improvement that can be directly compared with its cost of provision and other costs and benefits in order to decide whether or not it is a "good thing" in terms of its overall impact on social welfare (or in Broome's terms "value") at the margin. In purely practical terms, it is therefore no surprise that applied economists have tended to focus their energies on the empirical estimation of individual marginal rates of substitution of wealth for survival probability, leaving the selection of appropriate distributional weights to the relevant social decision-making agency. This is essentially the theoretical and ethical underpinning of the so-called "willingness to pay" (WTP)-based approach to the valuation of safety.

Clearly, if δp_i is set equal to $\frac{1}{n}$ for all n individual's whose utility appears in the social welfare function, then the mathematical expectation of the number of lives saved is precisely equal to one and the safety improvement is therefore referred to by economists as involving the saving of one "statistical life". If, in addition, for the sake of simplicity we set the marginal social welfare of wealth, $a_i \frac{\partial u_i}{\partial w_i}$, equal to unity across all individuals, then from equation (4) it follows that the gain in social

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

welfare from saving one statistical life – or in an economist's terms the "value of preventing one statistical fatality" (VPF) – is given simply by the population mean marginal rate of substitution of wealth for survival probability.

From an empirical point of view, these individual marginal rates of substitution have typically been estimated *either* by the "Revealed-Preference" approach in which the rate at which individuals actually do trade off wealth or income against survival probability (as in, say, labour markets) is the focus *or* by the so-called "Stated-Preference" approach which essentially involves asking respondents in a representative sample survey how much they would each individually be willing to pay for small reductions in their own risk of premature death by a given cause.

3. SINGING QALY-WALY DOODLE ALL THE DAY

Another way in which applied economists have approached the problem of dealing with variations in health and safety is via the concept of so-called "Quality-Adjusted Life Years" or, as they have come to be known in the trade, QALYs. As Broome notes in *Weighing Lives*, estimation of the number of QALYs associated with a given (typically impaired) health state involves establishing the time, T, spent in normal health that a representative sample of the population would regard as being equivalent to (i.e., no better or worse than) spending, say, ten years in that impaired health state. Thus, if T turned out to be, say, eight years then x years in the impaired health state would be treated as yielding 0.8x QALYs.

As such, it is essential to appreciate that defined and estimated in this way, QALYs are every bit as much based on the expressed preferences of members of the public as are willingness to pay-based monetary values of safety estimated using the "Stated-Preference" approach and hence every bit as much vulnerable to misperception and misrepresentation biases.

In principle, scarce healthcare resources are then allocated within the overall health budget on the basis of cost-effectiveness analysis, where the effectiveness of a medical treatment is measured by the QALYs gained as a result of the treatment, with the overall healthcare budget being allocated (again, in principle) with a view to minimizing the cost per QALY gained or, equivalently, maximizing the number of QALYs gained within the given budget.

Broome's main reservation concerning this approach is that it focuses exclusively on the improvement in a patient's health status that results from a given medical treatment and effectively ignores other (possibly very important) determinants of their overall wellbeing, such as waiting time, distance travelled to receive treatment, information provided by the medical profession and so on. As a result, Broome suggests that QALYs

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

should be replaced by "Wellbeing – Adjusted Life Years" or WALYs, whose scope would be sufficiently comprehensive to incorporate *all* of those factors related to medical treatment that might reasonably be held to be of concern to patients. But of course one is then bound to ask why, under Broome's suggested approach to the measurement of healthcare benefits, attention should be restricted only to those determinants of well-being affected by medical treatment – surely to be in line with the general ethical message of *Weighing Lives*, Broome's WALYs should take account of *all of* the determinants of an individual well-being. But then we are back to the sixtyfour trillion to the power of n-dollar question of how on earth such overall well-being is to be measured in practice.

Not surprisingly, Broome's reservations concerning QALYs are, at least to some degree, shared by others operating in this field. In addition, there are arguably at least three other major problems associated with QALYs as they are currently used in the allocation of healthcare resources. First, no account is currently taken of the personal circumstances of an individual enjoying a given QALY-gain i.e., his/her age, number of dependents etc. Second, if the concept of a QALY is to be given a sound theoretical underpinning then this would appear to require that lifetime well-being or utility should display strong (or additive) separability on the time dimension and this would seem, to say the least, to sit somewhat uncomfortably with the notion that for many people there may be strong intertemporal interactive effects at work in determining overall lifetime well-being. For example, does everyone regard five good years followed by five bad years as being no better and no worse than five bad years followed by five good ones, or some other mix of good/bad alternation? Third, if a particular healthcare programme generates other non-health related benefits, how are these to be aggregated with QALY gains in order to conduct an overall cost–benefit assessment?

Happily, Broome's concerns about QALYs, as well as those just outlined, may not remain entirely unanswered as The National Coordinating Centre for Research Methodology, on behalf of NICE and the NHS, has recently commissioned a research project which is being carried out jointly by a team drawn from the Universities of Aberdeen, East Anglia and Newcastle – of which the author is a member – to investigate these and other issues and, *inter alia*, to attempt to place a willingness to pay-based money value on a QALY or, in terminology that John Broome might prefer, to establish the "worth of a WALY", though it should of course be appreciated that the concept of well-being that a team of mere mortal researchers will be able to deal with will not *remotely* approach the comprehensive generality of the ethical jewel required by the argument developed in *Weighing Lives.*

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

4. HEADMASTER'S REPORT

So, sitting in his hard-backed chair in the Headmaster's study, how does John Broome react to an applied economist's attempts to deal with allocative and regulatory decisions that can be expected to affect current and future members of the human race? The following are some of the key points made by Broome in *Weighing Lives* and earlier publications, together with responses from a naughty schoolboy who unashamedly subscribes to the applied economist's view.

- In *Weighing Lives,* p. 263, it is asserted that:
  "To make a good assessment of the value of a person's life requires some theory. Yet the willingness-to-pay approach applies no theory. For example, it makes no presumption that extending a person's life by thirty years is more valuable than extending it by one year."

  In response to this criticism I would simply refer John Broome to pages 80, 81 and 103–9 of my 1976 book and Chapter 3 of its 1989 successor in which it is explicitly and unequivocally assumed that individual well-being is, *ceteris paribus,* an increasing function of longevity. He might also care to cast his eye over a 2004 report *Valuation of Health Benefits Associated with Reductions in Air Pollution* published by DEFRA and co-authored by myself and several colleagues, the central argument of which is quite clearly founded on the premise that increased longevity is positively valued by members of the public. In short, I would argue that the willingness to pay approach has a substantial theoretical underpinning, though I would be the first to concede that, like all theoretical foundations, it is open to many potential questions and criticisms. While he and I appear to disagree about most other things in this area, Charles Blackorby seems to share this view when he describes my 1989 book as being "– a book filled with more formulae and things I would normally have called theory than the usual economic treatise" (Blackorby 1990: 764)
- Again, in *Weighing Lives* and elsewhere Broome argues that in matters of life and death, most people are not the best judges of their own well-being and have, at best, only partially formed and often badly behaved preferences as judged by idealized standards, so that there is arguably a case for other better-informed, more "rational" agents to judge on their behalf.

  While I would have no hesitation in agreeing that some sort of well-informed expert input is highly desirable in public sector and related decisions that affect the health and safety of members of the public, I would also argue strongly in favour of taking full account of the latters' preferences, given that some attempt has been made to ensure

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

that these preferences have been as fully informed as is practically feasible.

- By now there exists extensive empirical evidence indicating that responses to contingent valuation (cv) (or "stated preference") questions aimed at estimating individual willingness to pay for reductions in the risk of premature death or injury display serious "insensitivity to scope" effects – see for example Beattie *et al.* (1998). That is, when asked about personal willingness to pay for, say, a 1 in 100,000 reduction in the risk of death during the coming year and then, at a later stage in the same interview, their corresponding willingness to pay for a 3 in 100,000 reduction, up to 40% of respondents gave the *same* answer, yielding radically different estimates of their marginal rate of substitution of wealth for survival probability.

  In *Weighing Lives* Broome identifies this as one of the key deficiencies of the WTP approach and, though he does not say so explicitly, his message to the muddy-booted boys in their end-of-term report, were he required to write one, would be: "Must try harder."

  In fact, we *did* try harder – indeed in terms of the time spent tussling with this problem, we tried *very much* harder. Our first conclusion was that our original CV questions suffered from the deficiency that Broome refers to elsewhere as being essentially "unstructured" – see Broome (1999) – that is, they required respondents to make what may be for many of them an excessively demanding comparison between variation on two radically different dimensions i.e., wealth on the one hand and tiny variations in the risk of premature death on the other. In view of this, we decided to employ a rather more highly structured approach by breaking the wealth-risk of premature death trade-off into a number of conceptually more manageable steps. In particular, we first asked respondents about their willingness to pay for the *certainty* of a cure for a non-fatal injury involving no permanent disability and their corresponding willingness to accept compensation for the certainty of incurring the same injury. On the basis of these responses and generally plausible assumptions concerning underlying preferences and attitudes to risk, it is then possible to infer the respondent's marginal rate of substitution of wealth for risk of the non-fatal injury. The second stage in the estimation process then involved the use of so-called "risk-risk" questions of the "standard gamble" type to estimate the respondent's willingness to trade-off risk of the non-fatal injury against risk of premature death. Finally, the estimated MRS of wealth for risk of the non-fatal injury was "chained" to the response to the risk-risk question in order to infer the respondents' implied rate of trade-off of wealth for risk of premature death. Happily, as reported in Carthy *et al.* (1999), the results displayed rather fewer scope-type

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

and other internal inconsistencies than did the results of our earlier studies, so that we hope that next term's Headmaster's report will have a rather chirpier feel to it.

- Empirical evidence tends to confirm the intuitive expectation that individual willingness to pay for safety will be an increasing function of *ability* to pay i.e., income or wealth. As a result, willingness to pay-based values of safety for poorer areas within a given country can be expected to stand at a significant discount relative to corresponding values for richer areas. Similarly, WTP-based VPFs for the developed world will substantially exceed those for developing countries. On this issue, the view from the Headmaster's study seems to be firm and unequivocal: some sort of corrective therapy in the form of a caning with distributional weights is clearly called for and, as far as the Head is aware, no such weights have so far been applied to the findings of the feckless applied economics pupils currently cavorting in the academic playground. In short, it would appear that the willingness-to-pay approach will inevitably value the safety of the rich more highly than that of the poor and in Broome's view this is simply ethically unacceptable.

  As one of those threatened with a whacking from the Head, I plead innocence on the following grounds. In the first place, I believe that distributional weights *are* effectively being applied, given the way in which the findings of studies aimed at the empirical estimation of WTP-based values of safety are actually applied in practice by government departments and related agencies. Thus, rather than using a high WTP-based VPF in the appraisal of proposed road safety improvements in, say, Belgravia and a low value in Stepney, the UK Department of Transport (DoT) in fact applies a *single* figure to *all* areas in the UK, this figure being computed as an *unweighted* average estimated from a representative sample taken from the population as a whole. To all intents and purposes this is tantamount to the application of distributional weights that are *inversely related* to income in that if, as empirical evidence tends to suggest, willingness to pay for safety improvement is broadly proportional to income, then multiplying each individual's willingness to pay by an inverse-income distributional weight will produce a weighted willingness to pay figure *that is the same for all individuals*. Since it seems to be natural to "scale" the distributional weights so that the weight applied to a person of average income is equal to unity (with weights that are larger than one applied to those of below average income and less than one to those above average) – and since empirical evidence tends to confirm the natural expectation that those on average income will have roughly average willingness to pay for safety – then the DoT's approach appears to

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

be entirely justified from an ethical point of view. To be sure, he might well put away his cane and instead punish the recalcitrants with what many would regard as the rather more gruesome retribution of having to read Jones-Lee (2004) in which it is argued that (viewing the distributional weighting problem from a between- rather than within-country perspective) a purely altruistic developed country should, in fact, value the safety externalities that it generates for a developing country at the latter's own "low" VPF, rather than its own "high" figure and that to do otherwise (by applying, say, a uniform value) would be Pareto – suboptimal. But having ploughed their way through that less than wholly merry missive, the deviants could continue to plead their case by pointing out that the "differential valuation" result follows only if the rich country displays pure altruism and that if instead it takes a *paternalistically* altruistic position, imposing its own preferences on the developing country's citizens (rather than respecting the latter's wishes) then according to Jones-Lee (2004) an "equal valuation" result emerges. The reader will simply have to make up his or her own mind as to whether or not the Head should again get out his cane.

- Current decisions concerning health and safety inevitably have some effect upon who will and who will not be alive in future generations and so how the human race will and will not develop. Since willingness to pay-based values of health and safety effectively take no account of such effects, not surprisingly John Broome takes the view that they are therefore seriously myopic.

  My first response to this criticism is that inasmuch as current public sector and related decisions have identifiable effects on the health and safety of future generations, then so far as is possible a well-executed willingness to pay-based study *will* attempt to take account of such effects. This does, of course, require fairly heroic assumptions concerning the likely nature of the preferences of members of future generations, but as in most other aspects of application of the willingness to pay approach, one simply has to try to do one's best. But it has to be conceded that this does not really come to grips with what seems to me to be the central and most problematic philosophical issue raised by the "future generations" question, namely who will and will not come into being and will the future existence or non-existence of a particular person add or subtract from Broome's overall measure of well-being of the past, present, and future members of the human race? I certainly don't even have the beginning of an answer to this question and nor, I suspect does John Broome or, for that matter, anyone else who's opinion I respect. So for the time being, I guess we'll all just have to muddle through and set the question delicately on one side.

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

APPENDIX: AN UNRESOLVED PUZZLE

Shortly before the symposium – but after this paper was written – John Broome raised with me a fascinating point, to which I have given some thought but must admit that, as yet, have no conclusive answer.

As pointed out in the main text, the standard result that the principal component of the willingness to pay-based VPF is the arithmetic mean over the affected population of individual marginal rates of substitution, m_i, of wealth for survival probability is based, *inter alia*, on the assumption that the marginal social welfare of the i^{th} individual's wealth, $a_i \frac{\partial u_i}{\partial w_i}$, is constant across all individuals. This leads essentially to the "standard" cost–benefit result that the value of a safety improvement is the aggregate *unweighted* willingness to pay by those individuals who benefit from the improvement.

Prima Facie, as noted in the main text, it would appear that this approach to the valuation of safety is of questionable fairness in that, literally applied, it would entail that the VPF for a low per capita income subsection of society would be lower than that for a higher income group. Again, as already noted, advocates of the approach would, however, counter this argument by pointing out that in practice, public sector bodies such as the UK Department for Transport apply a *common* VPF to *all* subsections of society, where this VPF is computed on the basis of the mean of m_i estimated from a representative sample of the *whole* of society.

In response to this, John Broome has suggested that an arguably more appropriate way in which to ensure a fair allocation of resources to safety improvement would be to set the marginal social welfare of the i^{th} individual's survival probability, $a_i \frac{\partial u_i}{\partial p_i}$, - rather than the i^{th} individual's wealth, $a_i \frac{\partial u_i}{\partial w_i}$ – constant across all individuals. On the face of it this, I must admit, seems to me to make a lot of sense.

However, things turn out to be not quite so straightforward. Thus, while the "standard" approach produces the result that $VPF = \frac{1}{n} \sum_i m_i$ (i.e., the arithmetic mean of m_i), a little algebra indicates that under Broome's condition and assuming that the safety improvement is financed by a *constant per capita* tax levied on all affected individuals, $VPF = \frac{1}{\frac{1}{n} \sum_i \frac{1}{m_i}}$ (which is the so-called *harmonic* mean of m_i). Clearly, with m_i constant across all individuals, the two definitions produce *exactly the same* VPF. However, it is a well-established result that for a given sample of non-negative, non-constant observations, the arithmetic mean will *always exceed* the harmonic mean. (The original proof of this result is apparently due to Augustin-Louis Cauchy using a very elegant, non-standard, induction process – see Aigner and Ziegler 2004: ch 17.)

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

The first thing to note, therefore, is that with m_i differing across individuals then the Broome approach will *necessarily* produce a VPF that is *less than* aggregate willingness to pay for the prevention of a statistical fatality. Furthermore, under Broome's approach it would, in principle, be possible to drive the VPF as close to zero as one wished by having m_i for at least one individual sufficiently close to zero. Indeed, with just one m_i *actually equal to zero*, then it would necessarily follow that the Broome VPF would itself be driven to zero. Given that in past empirical studies that I and others have conducted, genuine (i.e., non-protest) zero willingness to pay for safety improvement responses *have* emerged, this is, to say the least, a somewhat disturbing conclusion.

Quo Vadimus? In seeking a solution to this puzzle, two things should probably be borne in mind. First, while constancy of the marginal social welfare of wealth – which is assumed to be the case in deriving the standard result for the VPF – can in fact be shown to be a *necessary condition* for constrained social welfare maximization, this condition follows *only if* the tax levied on each individual in order to finance a safety improvement *is allowed to differ* from one individual to another. In turn, as already noted, Broome's result requires that the tax levied *is held constant* across all individuals. In the real world, it seems highly unlikely that *either* of these conditions on individual taxes would be politically feasible. So, anyone sufficiently interested in resolving this puzzle may wish to explore the implications of more realistic taxation procedures than either I or John Broome have assumed to be applicable.

POSTSCRIPT

Since the above Appendix was written, colleagues and I have succeeded in proving that given John Broome's suggested requirement that the distributional weights, a_i, should be chosen so that $a_i \frac{\partial u_i}{\partial p_i}$ is constant across all individuals, then provided that:

- the safety improvement is financed by a *proportional* tax on wealth levied on all individuals in society, and
- the wealth elasticity of m_i is equal to unity (which is broadly in line with empirical evidence)

then *whichever* subgroup of society actually benefits from the safety improvement, the VPF should be set equal to the population mean of m_i.

A paper providing a proof of this result is currently under submission, but a copy of the manuscript can be made available on request.

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press

REFERENCES

Aigner, M. and G. M. Ziegler. 2004. *Proofs From the Book*. New York, Springer.

Beattie, J., J. Covey, P. Dolan, L. Hopkins, M. Jones-Lee, G. Loomes, N. Pidgeon, A. Robinson, and A. Spencer. 1998. On the contingent valuation of safety and the safety of contingent valuation: Part 1 – Caveat Investigator. *Journal of Risk and Uncertainty* 17: 5–25.

Blackorby, C. 1990. Economic policy in a second-best environment. *Canadian Journal of Economics* XX111:748–71.

Broome, J. 1999. *Ethics out of Economics*. Cambridge, Cambridge University Press.

Chilton, S., J. Covey, M. Jones-Lee, G. Loomes, and H. Metcalf. 2004. *Valuation of Health Benefits Associated with Reductions in Air Pollution*. London, Department of the Environment, Food and Rural Affairs.

Carthy, T., S. Chilton, J. Covey, L. Hopkins, M. Jones-Lee, G. Loomes, N. Pidgeon, and A. Spencer. 1999. On the contingent valuation of safety and the safety of contingent valuation: Part 2 – The CV/SG 'Chained' Approach. *Journal of Risk and Uncertainty* 17:187–213.

Harsanyi, J. 1955. Cardinal welfare, individualistic ethics and interpersonal comparisons of utility. *Journal of Political Economy* 63:309–21.

Jones-Lee, M. W. 1974. *The Value of Life: An Economic Analysis*. London, Martin Robertson and Chicago, University of Chicago Press.

Jones-Lee, M. W. 1989. *The Economics of Safety and Physical Risk*. Oxford, Basil Blackwell.

Jones-Lee, M. W. 2004. Valuing international safety externalities: does the 'golden rule' apply? *Journal of Risk and Uncertainty* 29:277–87.

https://doi.org/10.1017/S0266267107001563 Published online by Cambridge University Press